

12
CLAIMS

1. A method, comprising:

a link receiver providing a plurality of data credits to a link transmitter;

5 the link transmitter transmitting a packet to the link receiver on an ingress link;

diminishing the plurality of data credits as the packet is transmitted;

the link receiver storing the packet in a plurality of receiver buffers;

the link receiver transmitting the packet out of the plurality of receiver buffers on a
egress link;

10 placing the plurality of receiver buffers into a free buffer pool as the packet is
transmitting out of the plurality of receiver buffers, wherein the free buffer pool
corresponds to additional data credits; and

the link receiver transmitting a flow control packet to the link transmitter, wherein
the flow control packet comprises the additional data credits.

15

2. The method of claim 1, wherein the ingress link has an ingress link speed, and
the egress link has an egress link speed, wherein placing the plurality of receiver buffers
into the free buffer pool comprises:

if the egress link speed is less than the ingress link speed, placing the plurality of
20 receiver buffers in the free buffer pool after a portion of the packet has been transmitted
out of the plurality of receiver buffers, and wherein the portion of the packet is
proportional to a ratio of the egress link speed to the ingress link speed; and

if the egress link speed is one of greater than and equal to the ingress link speed,
placing the plurality of receiver buffers into the free buffer pool when the packet begins
25 transmitting out of the plurality of receiver buffers.

3. The method of claim 2, wherein the packet begins transmitting out of the
plurality of receiver buffers when one of the plurality of receiver buffers is empty.

30 4. The method of claim 2, wherein the portion of the packet is substantially equal
to one minus the ratio of the egress link speed to the ingress link speed.

5. The method of claim 1, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link transmitter selects to which of the plurality of logical channels to allocate the additional data credits.

5 6. The method of claim 1, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link receiver selects to which of the plurality of logical channels to allocate the additional data credits.

7. The method of claim 1, wherein the link transmitter and the link receiver
10 operate in a switch fabric network.

8. The method of claim 7, wherein the switch fabric network is one of an Infiniband network and a Serial RapidIO network.

15 9. A method, comprising:
a link transmitter transmitting a packet to a link receiver on an ingress link;
diminishing a plurality of data credits at the link transmitter as the packet is
transmitted;
the link receiver storing the packet in a plurality of receiver buffers;
20 the link receiver transmitting the packet out of the plurality of receiver buffers on a
egress link; and
placing the plurality of receiver buffers into a free buffer pool when the packet
begins transmitting out of the plurality of receiver buffers, wherein the free buffer pool
corresponds to additional data credits.

25 10. The method of claim 9, wherein the packet begins transmitting out of the plurality of receiver buffers when one of the plurality of receiver buffers is empty.

11. The method of claim 9, further comprising wherein the link transmitter has a
30 plurality of logical channels, and wherein the link transmitter selects to which of the plurality of logical channels to allocate the additional data credits.

12. The method of claim 9, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link receiver selects to which of the plurality of logical channels to allocate the additional data credits.

5 13. A method, comprising:

a link transmitter transmitting a packet to a link receiver on an ingress link;
diminishing a plurality of data credits at the link transmitter as the packet is
transmitted;

the link receiver storing the packet in a plurality of receiver buffers;

10 the link receiver transmitting the packet out of the plurality of receiver buffers on a
egress link; and

placing the plurality of receiver buffers in a free buffer pool after a portion of the
packet has been transmitted out of the plurality of receiver buffers, wherein the portion of
the packet is proportional to a ratio of an egress link speed to an ingress link speed, and
15 wherein the free buffer pool corresponds to additional data credits.

14. The method of claim 13, wherein the portion of the packet is substantially
equal to one minus the ratio of the egress link speed to the ingress link speed.

20 15. The method of claim 13, further comprising wherein the link transmitter has a
plurality of logical channels, and wherein the link transmitter selects to which of the
plurality of logical channels to allocate the additional data credits.

25 16. The method of claim 13, further comprising wherein the link transmitter has a
plurality of logical channels, and wherein the link receiver selects to which of the plurality
of logical channels to allocate the additional data credits.

17. A switch, comprising:

a plurality of receiver buffers coupled to receive a packet from a link transmitter,
30 wherein the packet is stored in the plurality of receiver buffers, and wherein the switch
transmits the packet out of the plurality of receiver buffers;
a free buffer pool; and

a link receiver flow control algorithm, wherein the link receiver flow control algorithm places the plurality of receiver buffers into the free buffer pool as the packet is transmitting out of the plurality of receiver buffers.

5 18. The switch of claim 17, wherein the switch is coupled to receive the packet on an ingress link having an ingress link speed, and wherein the switch is coupled to transmit the packet on an egress link having an egress link speed, wherein placing the plurality of receiver buffers into the free buffer pool comprises:

 if the egress link speed is less than the ingress link speed, the plurality of receiver
10 buffers are placed in the free buffer pool after a portion of the packet has been transmitted out of the plurality of receiver buffers, and wherein the portion of the packet is proportional to a ratio of the egress link speed to the ingress link speed; and

 if the egress link speed is one of greater than and equal to the ingress link speed,
the plurality of receiver buffers are placed into the free buffer pool when the packet begins
15 transmitting out of the plurality of receiver buffers.

 19. The switch of claim 18, wherein the packet begins transmitting out of the plurality of receiver buffers when one of the plurality of receiver buffers is empty.

20 20. The switch of claim 18, wherein the portion of the packet is equal to one minus the ratio of the egress link speed to the ingress link speed.

 21. The switch of claim 17, wherein the switch operates in a switch fabric network.

25

 22. The switch of claim 21, wherein the switch fabric network is one of an Infiniband network and a Serial RapidIO network.

 23. A computer-readable medium containing computer instructions for instructing
30 a processor to perform a method of early buffer return, the instructions comprising:
a link transmitter transmitting a packet to a link receiver on an ingress link;
diminishing a plurality of data credits at the link transmitter as the packet is transmitted;

the link receiver storing the packet in a plurality of receiver buffers;

the link receiver transmitting the packet out of the plurality of receiver buffers on an egress link; and

5 placing the plurality of receiver buffers into a free buffer pool when the packet begins transmitting out of the plurality of receiver buffers, wherein the free buffer pool corresponds to additional data credits.

24. The computer-readable medium of claim 23, wherein the packet begins transmitting out of the plurality of receiver buffers when one of the plurality of receiver buffers is empty.

25. The computer-readable medium of claim 23, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link transmitter selects to which of the plurality of logical channels to allocate the additional data credits.

26. The computer-readable medium of claim 23, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link receiver selects to which of the plurality of logical channels to allocate the additional data credits.

20 27. A computer-readable medium containing computer instructions for instructing a processor to perform a method of early buffer return, the instructions comprising:

a link transmitter transmitting a packet to a link receiver on an ingress link;

diminishing a plurality of data credits at the link transmitter as the packet is transmitted;

25 the link receiver storing the packet in a plurality of receiver buffers;

the link receiver transmitting the packet out of the plurality of receiver buffers on an egress link; and

30 placing the plurality of receiver buffers in a free buffer pool after a portion of the packet has been transmitted out of the plurality of receiver buffers, wherein the portion of the packet is proportional to a ratio of an egress link speed to an ingress link speed, and wherein the free buffer pool corresponds to additional data credits.

28. The computer-readable medium of claim 27, wherein the portion of the packet is substantially equal to one minus the ratio of the egress link speed to the ingress link speed.

5 29. The computer-readable medium of claim 27, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link transmitter selects to which of the plurality of logical channels to allocate the additional data credits.

10 30. The computer-readable medium of claim 27, further comprising wherein the link transmitter has a plurality of logical channels, and wherein the link receiver selects to which of the plurality of logical channels to allocate the additional data credits.